

Transforming Science through Advanced Cyberinfrastructure

Manish Parashar
Office Director

Office of Advanced Cyberinfrastructure,

Directorate for Computer & Information Science & Engineering

National Science Foundation

ACCI Meeting
October 02, 2018

Outline

NSF/OAC Update Emerging
Science &
Role of Cl

Rethinking the CI Ecosystem

Conclusion

Outline

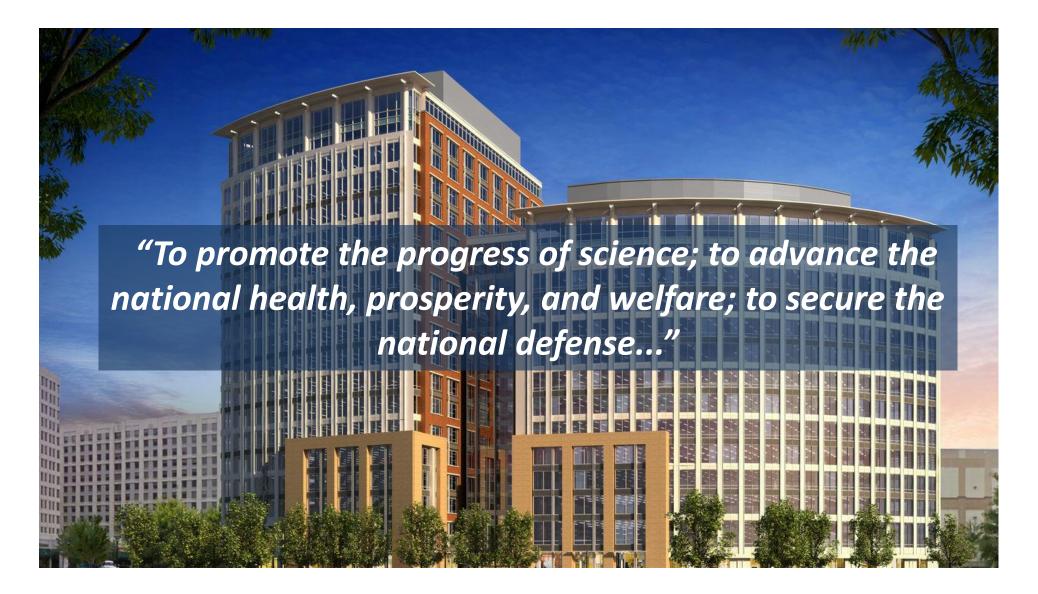
NSF/OAC Update Emerging
Science &
Role of Cl

Rethinking the CI Ecosystem

Conclusion



National Science Foundation's Mission





NSF Budget Update

FY 2018

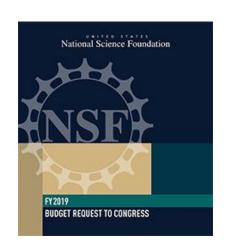
- FY18 Congressionally Appropriated Budget
 - +4%
 - +\$300M over FY17 budget!
- 2nd largest increase in NSF research budget increase in 15 years! (not counting ARRA)

FY 2019 (Request)



- Flat with respect to FY 17 Enacted
- CISE: \$925.4 million
 - -1.1% from FY 17 Enacted
- Big Ideas
 - Research Ideas: \$30 million each
 - Process Ideas: Midscale infrastructure: \$60 million
- Convergence Accelerators
 - \$60 million
 - HDR, FW-HTF: \$30 million each
- Continuing Resolution through December 07, 2018.





The NSF Big Ideas

RESEARCH IDEAS

HARNESSING THE DATA REVOLUTION

Harnessing Data for 21st Century Science and **Engineering**

Work at the **Human-**Technology Frontier: Shaping the **Future**



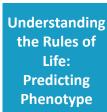






Quantum Leap: **Leading the** Next Quantum Revolution







" ... bold questions that will drive NSF's long-term research agenda -- questions that will ensure future generations continue to reap the benefits of fundamental S&E research."

PROCESS IDEAS

Mid-scale Research Infrastructure









Growing Convergence Research at NSF



NSF INCLUDES: Enhancing STEM through Diversity and Inclusion

Big Ideas => Big Cyberinfrastructure **Challenges & Opportunities**



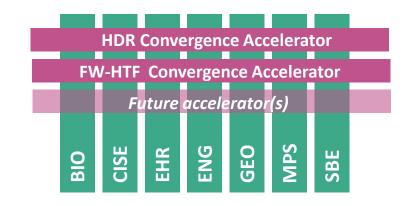
Convergence Accelerators

Accelerating Discovery through Convergence Research

Motivation: Changing nature of science research - research frontiers at intersection of existing disciplines

- Research: more intentionally managed, shorter timelines, milestones, deliverables, teams, partnerships. "DARPA-esque"
- Time-limited entities: accelerating impactful convergence research in areas of national importance
- Innovating in organizational structure: separate (from directorates) in leadership, budget, and programmatics
 - aligned with, relying on, foundational disciplinary research







Harnessing the Data Revolution (HDR)



HARNESSING THE DATA REVOLUTION

FUNDAMENTAL RESEARCH GOOD MACHINE

CYBERSECURITY DOMAIN

GOOD MACHINE

CYBERSECURITY DOMAIN

GOOD MACHINE

RESEARCH

GOOD MACHINE

VISUALIZATION

CYBERINFRASTRUCTURE

OF THINGS

MODELING

DATA

INTERNET

OF THINGS

MINING

HUMAN-DATA INTERFACE

"Engage NSF's research community in the pursuit of **fundamental research in data science and engineering**, the development of a
cohesive, federated, national-scale approach to **research data infrastructure**, and the development of **a 21st-century data-capable workforce**."

Research (across all NSF Directorates)

Theoretical foundations

Systems foundations data-centric algorithms systems

Data-intensive research

across all science & engineering

Educational pathways

Innovations grounded in an education-research-based framework





Advanced cyberinfrastructure

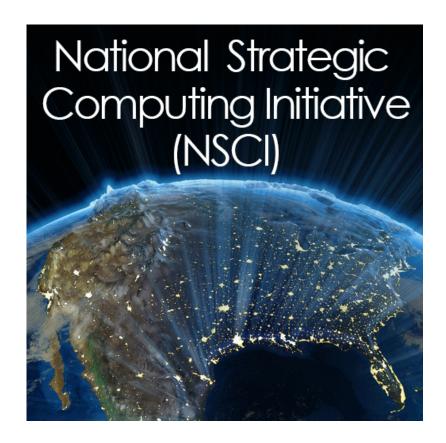
Accelerating data-intensive research.



National Strategic Computing Initiative (NSCI)

Maximizing benefits of HPC for scientific discovery and economic competitiveness

- Multi-agency effort to maximize the benefits of High Performance Computing (HPC) for scientific discovery and economic competitiveness
- NSF/CISE/OAC leads interagency effort to:
 - Increase coherence between technology base used for modeling/simulation and for data analytics;
 - Establish viable path forward for HPC systems in post-Moore's Law era; and
 - Increase capacity, capability, and sustainability of an enduring national HPC ecosystem



Includes CISE investments in the following programs: High Performance Computing (HPC), Scalable Parallelism in the Extreme (SPX), Quantum Leap Big Idea, and Cyberinfrastructure for Sustained Scientific Innovation – Software (CSSI)



EXECUTIVE OFFICE OF THE PRESIDENT WASHINGTON, D.C.



July 31, 2018

M-18-22

MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

FROM:

MICK MULVANEY

DIRECTOR, OFFICE OF MANAGEMENT AND BUDGET

MICHAEL KRATSIOS

DEPUTY ASSISTANT TO THE PRESIDENT

OFFICE OF SCIENCE AND TECHNOLOGY POLICY

NSF SUBJECT:

FY 2020 Administration Research and Development Budget Priorities



EXECUTIVE OFFICE OF THE PRESIDENT WASHINGTON, D.C.



American Leadership in Artificial Intelligence, Quantum Information Sciences, and Strategic Computing

Continued leadership in AI, quantum information science (QIS), and strategic computing is critically important to our national security and economic competitiveness. Advances in these areas promise opportunities for major scientific breakthroughs and are quickly transforming American life and industry. Agencies should invest in fundamental and applied AI research, including machine learning, autonomous systems, and applications at the human-technology frontier. Agencies should prioritize QIS R&D, which will build the technical and scientific base necessary to explore the next generation of **QIS** theory, devices, and applications. Agencies should prioritize investment in research and infrastructure to maintain U.S. leadership in strategic computing, from edge devices to high-performance computing, that accelerates delivery of lowpower, high performance devices; supports a national high-performance computing ecosystem; and explores novel pathways to advance computing in a post-Moore's Law era.

NSF Office of Advanced Cyberinfrastructure

Program Staff







Bill Miller Science Advisor

Computing

Data

Software

Networking & Cybersecurity

Learning & Workforce Development



Beth Plale*
Science
Advisor
Public Access



Bob Chadduck



Amy Walton



Vipin Chaudhary





Sushil Prasad



Alejandro Suarez Cooperative Agreements



Ed Walker



Stefan Robila





Kevin Thompson







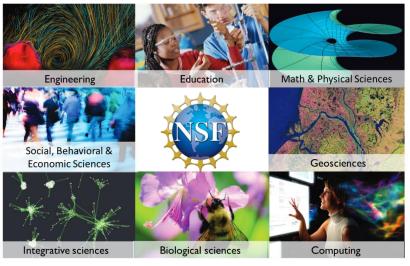
Join NSF/OAC: Multiple Program Officer openings

* IPA Appointment

CISE/OAC – Transforming the Frontiers of Science & Society

Foster a cyberinfrastructure ecosystem to transform computational- and data-intensive research across all of science and engineering

 Cyberinfrastructure Research & Research Cyberinfrastructure





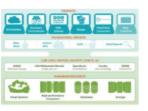
CI-Enabled
Instrumentation



Computing Resources



Data Infrastructure



Gateways, Hubs, and Services

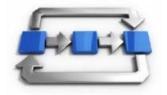




R&E Networks, Security Layers



Coordination & User support



Software and Workflow Systems



Pilots, Testbeds



People, organizations, and communities

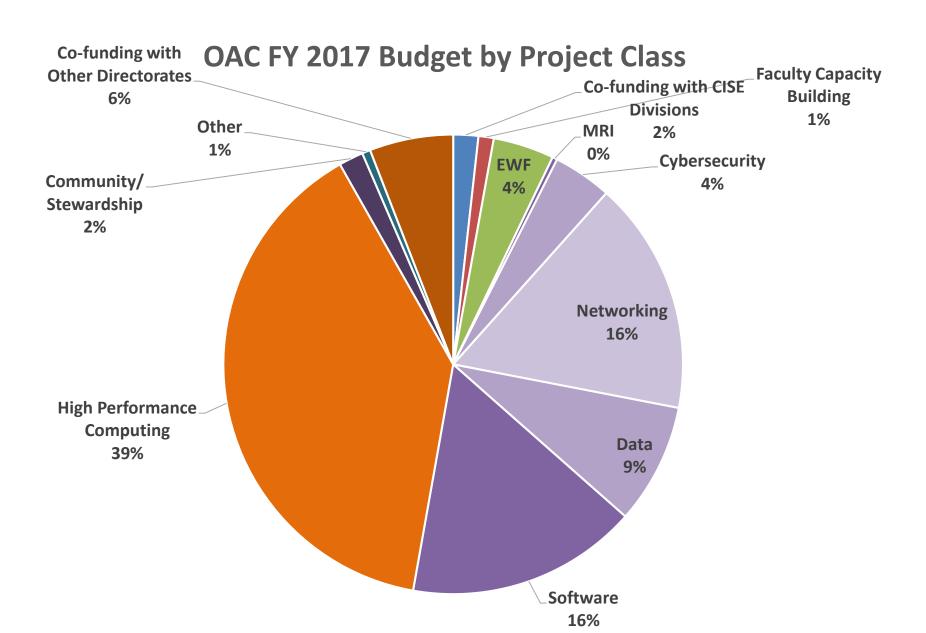


OAC by the Numbers: FY 2017





OAC Investments (2017)





Some Overarching OAC Principles

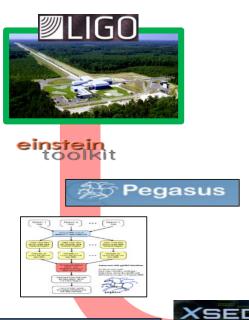
- **Science-driven**: Promotes science excellence, enabling fundamentally new scientific advances; benefits science and engineering communities beyond initial targets.
- Innovative: Emphasizes unique NSF contributions; builds the capability, capacity, and cohesiveness of a national CI ecosystem; considers both human and technical aspects of CI.
- Collaborative: Fosters partnerships and community development; actively engages CI experts, specialists and scientists working in concert with domain scientists who are CI users.
- Leveraged: Builds on existing, recognized capabilities.
- **Strategic**: Encourages measurement of progress and sharing of results through management plans and metrics.
- Sustainable: Provides benefits beyond the participants and the lifetime of the award.



Transforming Science through a Cyberinfrastructure Ecosystem: Dynamic discovery pathways at scale

CI enables Big Science

Gravitational wave detection enabled by NSF investments across the cyberinfrastructure ecosystem



globus

- **✓** Researcher access to sustained Advanced Computing resources
 - New intensive simulations of relativity and magnetohydrodynamics. Massive, parallel event searches and validation (100,000 models).
 - Advanced computing resources and services sponsored by NSF, DOE, and commercial cloud services.
- ✓ Interoperable Networking, Data Transfer, & Workflow Systems
 - Pegasus, HTCondor, Globus workflow and data transfer management
 - NSF funded 100 Gbps upgrades enabled huge throughput gains.
- ✓ Software Infrastructure
 - Computational science advances embodied in Software Infrastructure, for simulations, visualizations, workflows and data flows















NSF programs: Data Building Blocks (DIBBs), Software Infrastructure (SI²), Campus Cyberinfrastructure Network Infrastructure and Engineering (CC*NIE, DNI), and others. OSG and Pegasus are also supported by the Department of Energy.

Outline

NSF/OAC Update Emerging
Science &
Role of Cl

Rethinking the CI Ecosystem

Conclusion



Evolving Science, CI Landscapes

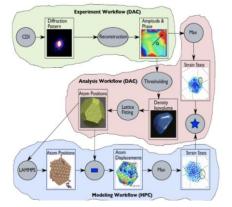
Evolving Science/Engineering Landscape

- Large scales, high-resolution, multi-scale, multiphysics simulations / Complex, dynamic workflows
- Emerging data-driven (ML-based) models
- Streaming data from observatories, instruments
- Rapidly growing "long-tail", "gateway" jobs; "small" jobs dominating; increasing use of clouds
- Heightened emphasis on robust results (transparency, credibility, correctness, security, ...)

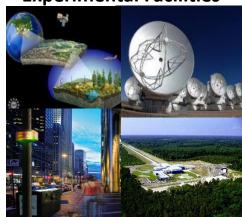
Evolving Technology Landscape

- Extreme scales / pervasive computing and data
- Diverse / disruptive technologies increasing
- Role of (non-traditional) software in taming complexity
- High throughput/low-latency networks
- Novel paradigms / Increasing role of clouds, multiclouds, hybrid environments / Growing capabilities & capacities at the edges

End-to-end Workflows



Instrument, Observatories, Experimental Facilities

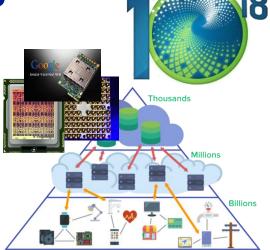


















Cyberinfrastructure is central to NSF's Large Facilities





Research success depends on robust, reliable, and highly connective cyberinfrastructure

Evolving Science, CI Landscapes

The cyberinfrastructure ecosystem must evolve to provide new levels of agility, interoperability, collaborations & reuse

How do we catalyze a **Cyberinfrastructure Ecosystem** to transform science?



Outline

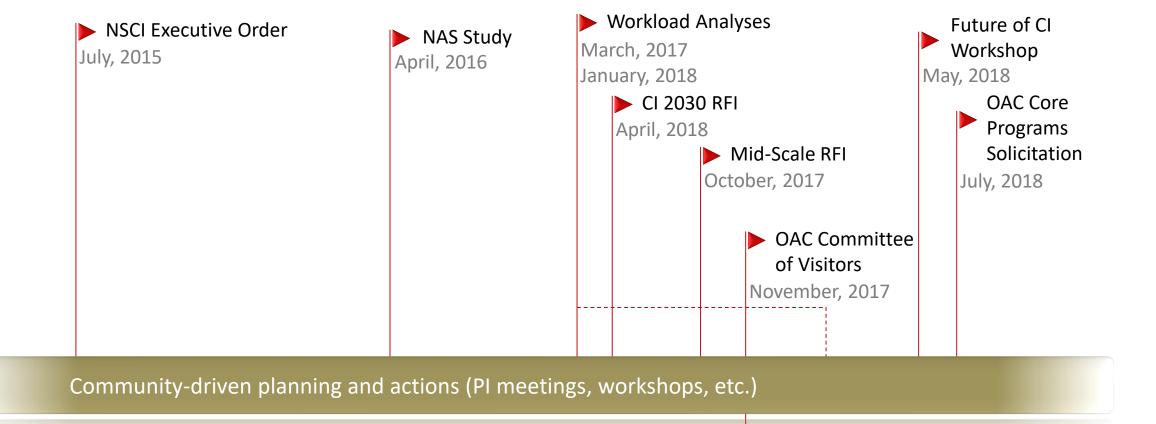
NSF/OAC Update Emerging
Science &
Role of Cl

Rethinking the CI Ecosystem

Conclusion



Rethinking CI: Building on Community Input



2017 -

Key Drivers

2015 -

Changing application landscape & workload profile

2016

Changing technology, services landscape

2018

- Increasing availability of (exp., obs.) data
- Growing role of ML, data-driven approaches

2019 →

Building on Community Input: Results of NSF CI 2030 Request for Information

Common needs expressed across science and engineering domains:

- Advanced computing. Growing need for <u>on-demand computing</u> for steering large simulations, rapid data processing, experiments; comparing simulations and observation.
- Data Science and management. <u>Big Data</u> and <u>Machine Learning</u>. Automated mining, analytics, visualization, provenance. Discoverability, accessibility, and reproducibility.
- Multi-source streaming data. Processing and integrating data from the <u>Internet of Things</u> (IOT) and cyber-physical systems at human, community, urban, and ecosystems scales.
- Secure access, dynamic and high bandwidth workflows. Technologies & approaches that scale with performance demands; storage, identity management, cybersecurity.
- Software. Porting, accelerating, validating algorithms and community codes. Software quality, reliability, validity, practices.
- Training and workforce development. For researchers and computing professionals, diversity and inclusion. CS/CI experts who collaborate closely with domain researchers.



All responses posted on CI 2030 Website: www.nsf.gov/cise/oac/ci2030/

Goal: Understand the nature and composition of the NSF-supported computational ecosystem in light of changing application requirements and resources and technology landscapes

- Looking forward, what are the essential components of, and relationships contributing to the NSF-supported computational ecosystem? What are the key application drivers (and usage modes) that motivate these components?
- What are the appropriate models for, and relationships contributing to resource management, operation and delivery?
- What are the opportunities, usage modes, and use cases for using future cyberinfrastructure components in an integrated and holistic manner?
- What the opportunities and use cases for integration with complementary investments by other agencies in the US and internationally?

https://uiowa.edu/nsfcyberinfrastructure/

Key Observations: The workshop highlighted the shifts in cyberinfrastructure needs and expectations, with a broad base of applications and resource needs that are much richer and more diverse.

- The demand for CI capabilities is dramatically expanding and diversifying;
 exceeds what NSF has currently provisioned.
- The high end (big data and big compute) and the long tail (laboratory and campus) of computationally-driven research are no longer distinguishable and deeply intertwined.
- Workforce training and availability remain challenges in campus
 cyberinfrastructure sustainability and support.

- Machine learning and big data analytics are now an essential part of the scientific discovery process, complementary to and increasingly integrated with computational modeling.
- Disciplinary and cross-disciplinary collaborations require converged instrument, data management, and computational capabilities.
- Incentives are needed for university investment to complement NSF's and to increase streamlined access to both, as well as commercial cloud services. NSF should not replicate what is well-supported by other agencies.
- Under current NSF budgets, there are neither adequate resources to meet the demand for midrange, campus hardware, services, and training nor to provision high-end hardware, services, and support comparable to those deployed by the EU, Japan, or China, or the Department of Energy. This should incorporate NSF's unique role in broadening participation.

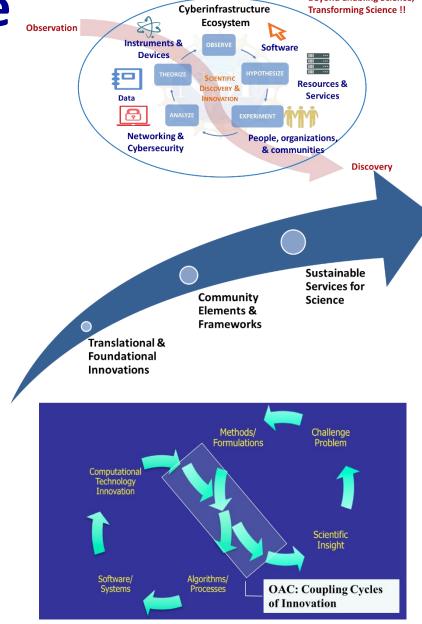
Early Recommendations:

- Delineate solicitations and investments in cyberinfrastructure innovation versus cyberinfrastructure operations, while recognizing the former ultimately informs the latter
- Consider a funding model that requires collaborations to drive interoperability and sustainability across scientific instruments and domains
- Explore new and creative kinds of partnerships public-private and interagency as necessary to sustain national research competitiveness and NSF leadership
- Develop a clear, long-term strategy, derived from principles that are clearly articulated and understood
 - Make difficult, strategic choices rather than investing in too many things, should focus limited resources on those things only it can do best



Rethinking CI: Realizing a CI Ecosystem to Transform Science Cyberinfrastructure Ecosystem

- Realize a holistic and integrated and secure cyberinfrastructure ecosystem aimed at transforming science
- Support the translational research continuum, from catalyzing core innovations, through fostering the community tools and frameworks, and enabling sustainable cyberinfrastructure services
- Work closely with science and engineering communities, and other stakeholders to tightly couple the cycles of discovery and innovation



Rethinking CI: Realizing a CI Ecosystem to Transform Science Cyberinfrastructure Ecosystem

- Implications on Programs and Solicitations
 - Prioritize/focus investments to maximize impact
- Define linkages / coordinate across the traditional OAC clusters and solicitations
 - Invest in CI research; support the translational arc
 - Leverage/complement investments by campuses, industry (e.g., Cloud services), other agencies
 - Engage with the stakeholders, community



OAC Core Research Program NSF 18-567; Due 11/15/18

CISE/OAC – Transforming the Frontiers of Science & Society

Computing

Advanced resources and services at all scales – MRI (clusters); Innovative HPC; Leadership Class; XSEDE coordination and user services; Research

Data

Data Building Blocks (DIBBS)
Program

Cyberinfrastructure for Sustained Scientific Innovation (CSSI)

Software

Software Infrastructure for Sustained Innovation (SI2)

Networking & Cybersecurity

Campus Cyberinfrastructure (CC*), International Research Network Connections (IRNC), Cybersecurity Innovation for CI (CICI)

Learning & Workforce Dev.

Training-based Workforce Development for Advanced Cyberinfrastructure (CyberTraining), CAREER, CRII



Cyberinfrastructure for Emerging Science and Engineering Research (CESER), Public Access, Big Data Hubs



OAC Core Research Program SOLICITATION NSF 18-567

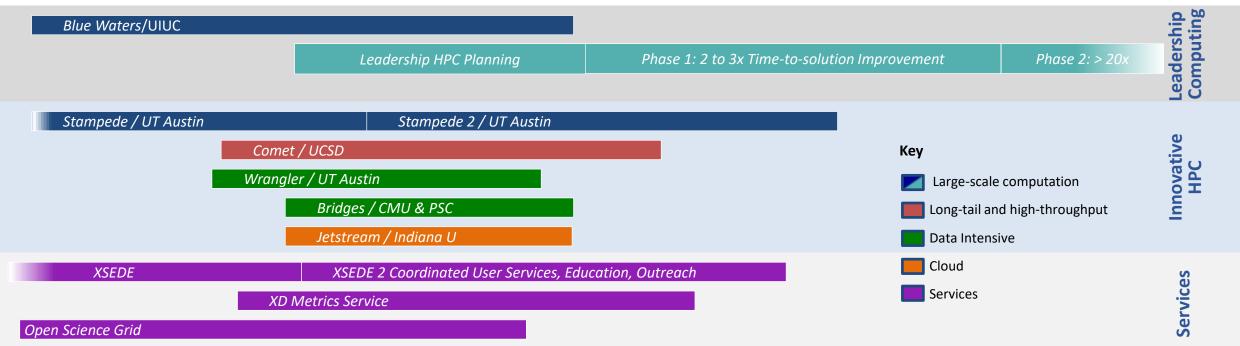


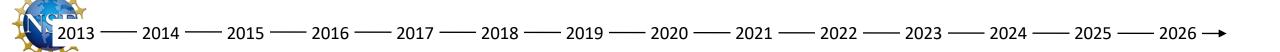
- *Translational research* (spanning design to practice) in all aspects of advanced cyberinfrastructure (CI) to transform science and engineering research
 - Multi-disciplinary, extreme-scale, driven by science and engineering research, end-to-end, or deployable as robust research CI
- Research Areas
 - Architecture and middleware for extreme-scale systems
 - Scalable Algorithms and Applications
 - Advanced Cyberinfrastructure Ecosystem
- Research Communities: Multiple disciplinary areas supported spanning Computer as well as Computational and Data-driven Science and Engineering

- Part of CISE's coordinated core program solicitations
 - Only Small proposals in FY'19
 - Funding amount \$7.5M
 - Max \$500K/award
- Proposals due Nov 15, 2018
 - Pl's strongly encouraged to send 1-page project summary for further guidance.
 - Webinar in July/Aug



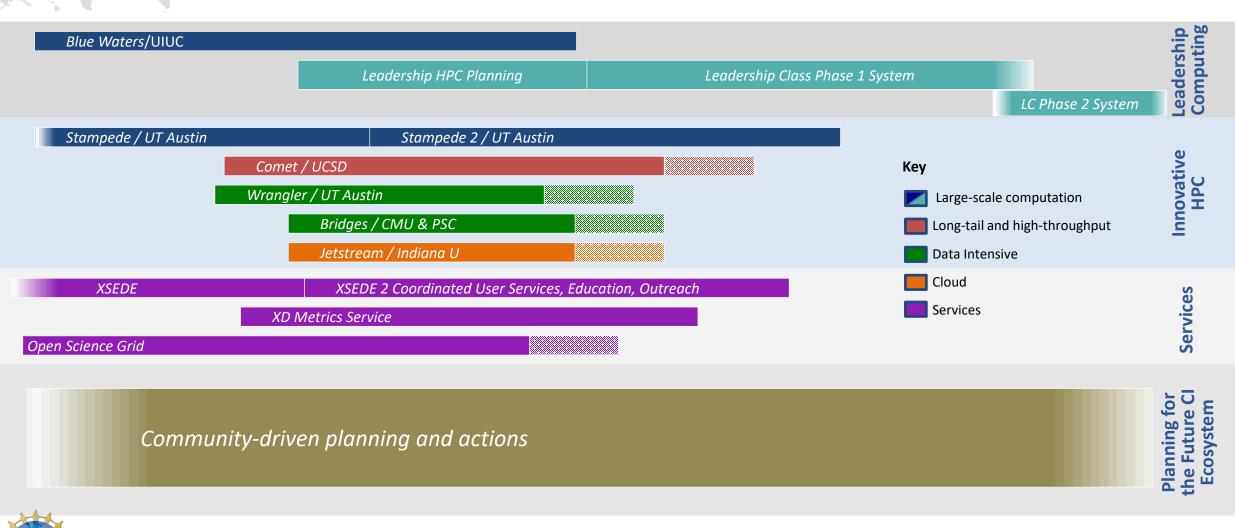
OAC-Funded Computing Ecosystem







OAC-Funded Computing Ecosystem



2014 — 2015 — 2016 — 2017 — 2018 — 2019 — 2020 — 2021 — 2022 — 2023 — 2024 — 2025 — 2026 —

Rethinking CI – Computational Infrastructure*

* Not mutually exclusive.

- Highly specialized instruments
- Extreme-scale capabilities
- NSF 17-558

Leadership Class Systems





NSF 17-558: Towards a Leadership-Class Computing Facility – Phase 1

Resolution

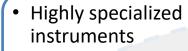
RESOLVED, that the National Science Board authorizes the Director at her discretion to make an award, OAC-1818253, to the Texas Advanced Computing Center (TACC) at the University of Texas at Austin for the acquisition of the system described in proposal "Computation for the Endless Frontier," in an amount not to exceed \$60,000,000 for a period of 60 months. Pending appropriate approval associated with NSF MREFC policies, an additional amount not to exceed \$8 million may be made available to TACC in the form of supplemental funding to this award to advance the design of the Phase 2 leadership-class system.



- To be deployed 2019
 - Primary compute: DellEMC and Intel
 - 35-40 PetaFlops Peak Performance
 - Interconnect: Mellanox HDR & HDR-100
 - Fat Tree topology, 200Gb/s bet. switches.
 - Storage: DataDirect Networks
 - 50+ PB disk, 3PB Flash, 1.5TB/sec peak.
 - Single Precision Subsystem: Nvidia
 - Front end for data movers, workflows
- Potential 10x Phase 2 system.

Rethinking CI - Computational Infrastructure*

* Not mutually exclusive.



- Extreme-scale capabilities
- NSF 17-558

Leadership Class Systems

Federated

Resources

Innovative Capacity Capabilities Support smallmedium scale computations, data analytics, long-tail users, gateway computation

S. WIETNER

Cloud Services

Campus C

 Integrated distributed CI elements (capacity, capability, campus, network, storage, services, etc.) Exploratory
Pilots &
Prototypes

- Explore novel technologies, architectures, usage modes, etc.
- Path for transition to production important



Coordination & Measurement

Engaging with the Community

- OAC webinar series
 - Focus on the translational impact of advanced cyberinfrastructure and how cyberinfrastructure innovations have impacted science.
 - 3rd Thursday of each month at 2:00pm EDT
 - https://www.youtube.com/playlist?list=PL0ujJTaPsv3eXvi3PJqo CSIkzFtfcsNbW
- Newsletters, community notifications



Town Halls at PEARC'18, SC'18, ...

Outline

NSF/OAC Update Emerging
Science &
Role of Cl

Rethinking the CI Ecosystem

Conclusion



Some Personal Observations

- Unprecedented opportunities for insights, discoveries: CI poised for transformative impact
- Compute and data are essential to all areas of science, engineering: easy/pervasive/cost-effective access to CI is critical
- Dramatic changes in application, method, technology, CI solution landscapes: an inflection point in cyberinfrastructure
- Growing demands across the spectrum; flat budgets: how to maximize the impact of NSF investments? How to leverage other investments (including industry)?
- CI ecosystem is a national asset: where can NSF investments add unique value as part of a larger CI ecosystem?

OAC's Big Questions for ACCI

Research Agenda:

- What constitutes CI research and what is the OAC research agenda?
- How does OAC foster/nurture its research community to ensure innovations in translational computer and computation science/engineering and CI?

Reproducibility:

 What is the unique role of CI and OAC in the broader conversation on reproducibility and repeatability of research results?

Integrated CI Ecosystem:

 How does OAC evolve its priorities and programs to address the growing need for an integrated CI ecosystem?

Sustainability:

- What is OACs role in sustainability of resources and services including compute, data/storage, networking, software, educational?
 - What are appropriate models for sustainability?
 - How does OAC decide what should be sustained?
- Difference between making resources/services sustainable and sustaining resources/services? Should OAC focus on the former and work with the community on the latter?
- How does OAC work with other directorates towards sustainability of specific science?
 Are there core CI services and infrastructure that is critical for all sciences?



OAC's Big Questions for ACCI

Strategic Investments:

- What is OAC's role in a rapidly changing application and technology/service landscape? Emerging applications span extreme scale, end-toend, computing and data intensive, data-driven, etc., and the emerging technology/service landscape includes accelerators, deep memory hierarchies, cloud/edge services, software defined networking, etc.
- How does OAC prepare for technology disruptions (e.g., beyond Moore's law, quantum, etc.)?
- How can OAC partner with industry? Are there specific areas of partnerships that are more productive?

Diversity and Broadening Participation:

 How can OAC effectively address broadening participation/diversity to support underrepresented groups and institutions without essential

- cyberinfrastructure resources?
- Are there unique challenges for OAC and how does OAC address them?

LWD:

- What are the research, education and reward structures, communities and publications, and career paths for current and future OAC research communities?
- What are the relevant conferences and journals and what is their perceived quality?
- What are mechanisms to increase engagement of Computer Science faculty in CI?

Communication:

 What are the most impactful dissemination mechanisms for communicating with the community and the broader public?



Conclusion

- Science and society are being transformed by compute and data

 a connected, robust and secure cyberinfrastructure ecosystem is essential
- Rapidly changing application requirements; resource and technology landscapes
 - Our cyberinfrastructure ecosystem must evolve in response
- Lets build a cyberinfrastructure ecosystem aimed at transforming science



Join the conversation

- OAC Webinar Series
 - 3rd Thursday @ 2PM ET
- OAC Newsletter
- Follow us on Twitter @NSF_CISE

Stay informed

- Join the OAC, CISE Mailing Lists
 - Learn about NSF events, programs, webinars, etc.
 - Send email to:
 - oac-announce-subscribe-request@listserv.nsf.gov
 - <u>cise-announce-subscribe-request@listserv.nsf.gov</u>

Get involved

- Reviews proposals, serve on panels
- Visit NSF, get to know your programs and Program Officers
- Participate in NSF workshops and visioning activities
- Join NSF: serve as Program Officer,
 Division Director, or Science Advisor

NSF Office of Advanced Cyberinfrastructure (OAC) Newsletter



"Make no little plans; They have no magic to stir men's blood ..."

Daniel H. Burnham, Architect and City Planner Extraordinaire, 1907.

"If you want to travel fast, travel alone; if you want to travel far, travel together"



Manish Parashar

Office Director, Office of Advanced Cyberinfrastructure

Email: mparasha@nsf.gov



To subscribe to the OAC Announce Mailing List Send an email to: OAC-ANNOUNCE-subscribe-request@listserv.nsf.gov



African Proverb.